

Fault Tolerance Requirements of Tactical Information Management Systems

Jeffrey Cleveland and Joseph Loyall
Raytheon BBN Technologies
Cambridge, MA, USA

James Hanna
Air Force Research Laboratory
Rome, NY, USA

Abstract— Information Management (IM) services provide a powerful capability for military operations, enabling managed information exchange based on the characteristics of the information that is needed and the information that is available, rather than on explicit knowledge of the information consumers, producers, and repositories. To be usable in tactical environments and mission critical operations, IM services need to be resilient to faults and failures, which can be due to many factors, including design or implementation flaws, misconfiguration, corruption, hardware or infrastructure failure, resource intermittency or contention, or hostile actions. This paper presents a reference model for representing the performance and fault tolerance requirements of IM services in tactical operations. A Joint Close Air Support operation is described using this representation and the viability of canonical fault tolerance techniques are examined for a given deployment.

Index Terms— Fault Tolerance, Information Management Systems, Military Operational Scenarios, System Requirements.

I. INTRODUCTION

Information Management (IM) services provide a powerful capability for military operations, enabling managed information exchange between tactical warfighters, command and control centers, surveillance and intelligence assets, and aircraft in service of ongoing missions. IM services, based on a publish-subscribe-query model, support an information-centric organization of a distributed system, rather than the traditional interface-centric organization. In an information-centric system, information providers and consumers are decoupled from one another. IM services manage information exchange based on the characteristics of the information that is needed and the information that is available, rather than on explicit knowledge of the information consumers, producers, and repositories.

The core concept of IM is active information management in which clients are information publishers and consumers that communicate with other clients via shared IM services, including publication, discovery, brokering, archiving, and querying [4], [6]. Sensors (such as those on manned or unmanned vehicles) and other information producers (such as tactical warfighters or information analysts) publish information. Consumers make requests for future information through subscriptions or for past information through queries.

To be usable in tactical environments and mission critical operations, IM services need to be survivable, i.e., resilient and adaptive to faults and failures with or without malicious

intelligence behind them. Faults and failures in a deployed system can be due to many factors, including design or implementation flaws, misconfiguration, corrupted processes or information, hardware or infrastructure failure, resource intermittency or contention, or hostile actions by adversaries.

Fault tolerance capabilities must address mission requirements based on their anticipated deployment patterns and use of IM services. Toward this end, we have produced a reference model for describing relevant properties of operational scenarios and deployments that could benefit from resilient IM services.

Our reference model consists of two main components, a scenario description and deployment descriptions. The scenario description includes elements that are mission specific, e.g., information sources and sinks, the properties of exchanged information, and how information is exchanged over time. Each deployment description describes a specific deployment of IM services to support a scenario including which actors are hosted on enterprise or embedded platforms, where IM services can be located (e.g., in centralized, distributed, or hybrid configurations), and the interactions between services and clients (including the potential communications available).

To derive fault tolerance requirements, we first use the scenario descriptions to define operational scenario requirements, which we then apply to each of the described deployment descriptions. We illustrate our reference model with a Joint Close Air Support (JCAS) scenario based upon exercises and demonstrations in which we have been involved and available documentation of military operations, doctrine, and guidance.

II. EXAMPLE SCENARIO

Throughout this paper we will apply our reference model to an example JCAS operational scenario. Close Air Support (CAS) is defined by [10] as “air action by fixed-wing and rotary-wing aircraft against hostile targets that are in close proximity to friendly forces and requires detailed integration of each air mission with the fire and movement of those forces.”

The JCAS scenario is derived from several documents, including accounts of previous operations such as those described in [8], [24], [26], and other studies of military operations such as [18] and [19]. Another resource that we draw upon is the Joint Force doctrine for various scenarios including [10], [11], and [12]. We also draw upon our prior experience in several live-flight and live-fire operational exercises [5], [16], [17], [22].

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Fault Tolerance Requirements of Tactical Information Management Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Raytheon BBN Technologies,10 Moulton Street,Cambridge,MA,02138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Military Communications Conference (MILCOM), Orlando, Florida, October 29-November 1, 2012.					
14. ABSTRACT Information Management (IM) services provide a powerful capability for military operations, enabling managed information exchange based on the characteristics of the information that is needed and the information that is available, rather than on explicit knowledge of the information consumers, producers, and repositories. To be usable in tactical environments and mission critical operations, IM services need to be resilient to faults and failures, which can be due to many factors, including design or implementation flaws, misconfiguration, corruption, hardware or infrastructure failure, resource intermittency or contention, or hostile actions. This paper presents a reference model for representing the performance and fault tolerance requirements of IM services in tactical operations. A Joint Close Air Support operation is described using this representation and the viability of canonical fault tolerance techniques are examined for a given deployment.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

In the JCAS scenario, a Joint Tactical Air Controller (JTAC) identifies a target, requests air support, and then transmits target data to a strike aircraft fulfilling the role of a CAS Aircraft (AC). The CAS AC then performs the requested strike. IM Services are utilized to provide situational awareness to all parties involved, minimizing the chance of fratricide and collateral damage while maximizing the effectiveness of the strike. Additionally, IM services are utilized for forwarding the strike request, target data, and coordination (including a potential abort command) between the JTAC, Command and Control (C2) entities, and the CAS AC.

III. OPERATIONAL SCENARIO DESCRIPTIONS

Each operational scenario description includes the mission components that are necessary for the assessment of deployment patterns and fault tolerance requirements. The descriptions include details such as the mission participants, the type of information exchanged, size and frequency of data transfers, available computing resources, and so forth. These descriptions are not intended to serve as exact documentation of a specific exercise or past mission, but rather as general use cases to consider while developing and assessing IM systems.

A. Information Data Profiles

The first piece of the scenario information flow descriptions are data profiles. These profiles outline attributes of the information passed during the scenario such as quantity, frequency, size, and time/reliability requirements. Such profiles are an important part of the IM requirements for each scenario and are vital when determining appropriate fault tolerant deployments and requirements. The following categories of information are at the core of the JCAS scenario.

Force Track / Location Data. In several scenarios, different types of location data are utilized, e.g., target locations that convey the position of something at a given time.

Blue Force Tracks (BFTs) convey the current position of friendly forces. BFT may be automatically published and should include at least GPS position, timestamp, and unit id. The Troops In Contact (TIC) Report is a specific type of BFT which indicates friendly forces in contact with hostile forces.

Red Force Tracks (RFTs) convey the current (estimated) position of enemy forces. RFTs are generally pushed out by friendly forces based on estimates or readings of ground sensors (e.g., a ground trip sensor may publish a RFT when it is triggered), and as such are often not updated as frequently (and are not as precise) as BFT. While how to best represent RFTs is actively researched [3], this data should at a minimum include positional data and a timestamp.

White Force Tracks are similar to RFTs in form and convey the current (estimated) position of neutral/unknown individuals.

Vehicular Track and Readiness Data. Vehicular track data can take a wide range of forms describing attributes of vehicle tracks, e.g., classification, vehicular type, and kinematic data. Depending on the type of vehicle transmitting the track and readiness data, this may also include attributes such as sensor status, weapons inventory, or current tasking information.

Image/Video Data. Transmitting visual data is a central use of IM systems in tactical settings. Videos and still images may be resized or re-encoded in transit to improve quality of service

(QoS). The size of the original image depends on the sensor and the size of the delivered image should depend on the receiving device, e.g., a massive gigapixel image can be scaled down before it is transmitted. A typical 640x480 JPEG encoded image is ~50 KB.

As with still images, the size of video files depends on the image resolution and encoding method, as well as frame rate. For example, at 30 frames per second, standard definition MPEG-2 will generally be in the 2-5 MB/second range.

Two common variations of standard visual data include annotated images/videos and georectified images/video. Annotated images are generally the same quality and resolution as the original, and the size should be approximately the same with some additional metadata describing the annotation. Georectified visual data is commonly used for directing weapon strikes. The georectification process involves identifying multiple tie points which are used to assign precision coordinate data to each pixel. This coordinate data can drastically increase the size of the original image. For instance, each pixel may go from being represented as a 16 bit value to being represented as a 128 bit value. The size of this target data will depend on the size of the image being used, e.g., a 640x480 JPEG may go from being ~50 KB to ~5 MB once it is georectified. Due to the precision required in weapon strikes, the quality of georectified data should not be reduced if possible.

Air Support Requests. Air support requests are a group of information type conveying the need for some form of air support. The 9-Line CAS Briefing, utilized in the example JCAS scenario, contains information relayed from ground forces that an aircraft needs to carry out CAS. This information is summarized by the following 9 lines: Initial position, Heading/Offset, Distance to Target, Target elevation, Target description, Target location, Type of mark on target or laser code, Location of friendlies, and Egress.

Similar data types include the 9-Line CAS Evac brief, which contains information necessary to request an air evacuation of troops, and the Joint Tactical Airstrike Request (JTAR), which contains the information necessary to request an air strike on a target including strike priority, the type of target, target location, time of strike, desired ordnance, call sign and frequency of unit requesting strike, and additional remarks.

Approval/Confirmation Data. Read backs are used to ensure that critical information, such as the details of an air strike, were conveyed properly. When communication is over a voice channel, the receiver reads back the information received so that each side can ensure it was heard properly. When information is being passed through an IM system this could entail the receiver retransmitting the original message in its entirety or transmitting a checksum of the information to be verified by the original publisher. Cleared Hot/Abort Message indicates that a strike is cleared to be carried out or should be aborted.

Tasking Related Data Types. Tasking objects are a notional data type used to describe where units should be moving. Depending on the type of unit, this may take many different forms. At a minimum, it is expected to contain the target location and potentially vectors or specific routes to take in order to reach that target.

B. Information Flow View

The information flow views describe how actors in a scenario share information. This is the core capability provided by IM services and as such is central to the scenario descriptions. We illustrate information flow in two ways: (1) as an operational view diagram showing the source and destination of information and (2) as a color coded grid where the y-axis indicates information type, the x-axis indicates actor, and each block in the grid indicates if that actor transmits and/or receives that type of information. Figs. 1 and 2 provide information flow diagrams for the example JCAS scenario.

C. Scenario Event Sequences

UML-like sequence diagrams are used to provide the sequence of events that occur as the scenario progresses, illustrating the temporal aspect of information exchanges within the scenario. Fig. 3 provides a partial sequence diagram for our example JCAS Scenario.

D. Scalable Dimensions of the Scenario

Our scenario descriptions also include ways in which scenarios are likely to increase in scale. Certain deployments of IM services may be sufficient for small scales. As the number of actors (such as ground forces) scale up, some of these deployment patterns may become infeasible. By identifying and documenting the aspects of a scenario most likely to increase, we hope to help match the scenario descriptions to potential system deployments.

There are multiple aspects of the JCAS scenario that are likely to scale. The number of JTACs is likely to increase based on the number of friendly forces in an area and the nature of ongoing operations. Additionally, the number of CAS AC is likely to increase with the level of activity inside an operational area. Similarly, the number of Forward Air Control Aircraft (FAC AC) is likely to increase based on the size of the operational area. Lastly, the number of Intelligence, Surveillance, and Reconnaissance (ISR) assets is likely to increase based on the size of the operational area.

E. Information Pipeline Categorization

We categorize groups of information sources and sinks

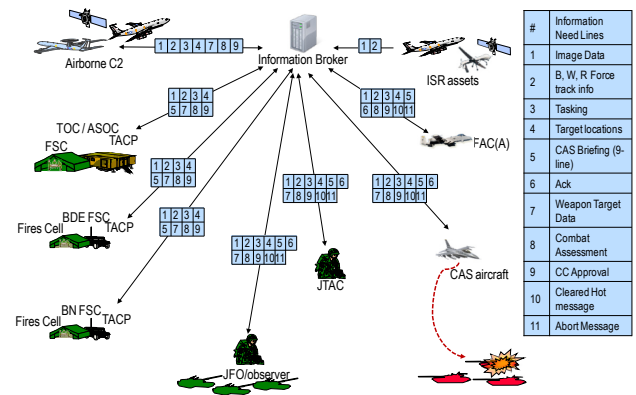


Fig. 1. Operational Information Flow View.

	Airborne C2	ISR Assets	ASOC	JTAC	JFO	FAC(A)	CAS aircraft	IMS
Video/image of target	T	T	T	T	T	T	T	T/R
Annotated image/video	T		T	R	T	T	T	T/R
Confirm RX				R		T	T	T/R
CAS Briefing (9 Line)				R	T	T	T	T/R
CAS Briefing "Read Back"				T	T	R	R	T/R
Final Target Data				T	T	R	R	T/R
Cleared hot message				R	T	T	T	T/R
Abort message				R	T/R	T	T	T/R
<div>Transmit (T)</div> <div>Receive (R)</div> <div>Transmit/Receive (T/R)</div>								

Fig. 2. Information Flow View.

based upon common attributes such as data type or purpose. This provides a useful abstraction when dealing with potentially long and complex data exchanges of an operational scenario.

Each scenario can be seen as consisting of information

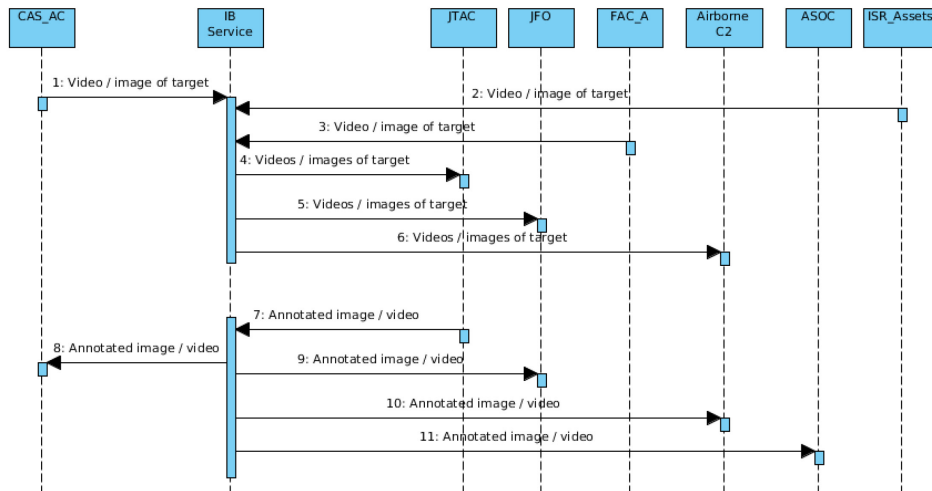


Fig. 3. A Partial JCAS Sequence Diagram.

pipelines that utilize the following IM services [6]:

- A *Submission* service to publish information, such as by a sensor.
- A *Brokering* service that matches published information to information requests, such as subscriptions.
- A *Dissemination* service that distributes information to participants that have requested it.

We identified the following five information pipelines in the JCAS scenario:

1. End-clients (both enterprise and embedded) exchanging situational awareness data including BFT, ISR data, and annotated ISR data.
2. JTAC requesting CAS by passing 9-line data to FAC AC, Joint Fires Observer (JFO), and CAS AC.
3. Tasking orders from C2 platforms to CAS AC.
4. Strike coordination between the JTAC and the CAS AC.
5. Hot/Abort message to the CAS AC which can originate from multiple actors in the field.

For each of these pipelines, a publisher sends information to a submission service where it is then forwarded to a brokering service, the proper subscribers are identified, and the information is then passed to a dissemination service which forwards it to each subscriber.

F. Information Pipeline Requirements

The information pipeline abstraction provides an adequate granularity for assigning information requirements without being bogged down by the many information sources and destinations in an operational scenario.

The JCAS pipeline requirements are shown in Table I. Each pipeline must provide a continuous flow of information between the actors involved. Pipeline 1 may allow for a reduction of information fidelity for non-targeting data. Pipeline 2, 3, 4 and 5 must not reduce information fidelity. The timeliness requirements for information pipelines 1, 2, 3, and 4 are less strict than for pipeline 5, which is responsible for carrying the final hot/abort message of the air strike and must be delivered within a time frame consistent with CAS doctrine.

Based upon the pipeline requirements, we define service requirements in terms of availability, timeliness, and integrity. In this scenario, to satisfy the strictest requirements of pipelines, the IM services must be continuously available, with immediate timeliness, and integrity guarantees.

TABLE I. JCAS PIPELINE REQUIREMENTS

	Reliability	Timeliness	Fidelity reduction
Pipeline 1	Unordered	Minutes	Acceptable
Pipeline 2	Unordered	Minutes	Unacceptable
Pipeline 3	Unordered	Seconds	Unacceptable
Pipeline 4	Ordered	Seconds	Unacceptable
Pipeline 5	Ordered	Immediate	Unacceptable

IV. DEPLOYMENT DESCRIPTIONS

The deployments that we describe consist of a physical communication network, actors and nodes attached to that network, and roles such as clients and services hosted on the

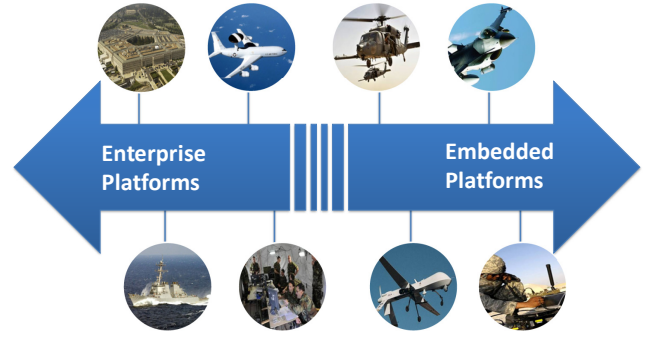


Fig. 4. Spectrum of Heterogeneous Tactical Platforms.

nodes. This section describes a means of documenting potential deployments to support defining fault tolerance requirements of a deployment for a specific operational scenario. We illustrate the deployment descriptions using the example of the JCAS operational scenario.

A. Platform Class Abstraction

Military networks can consist of various types of nodes which exhibit a high level of heterogeneity. These heterogeneous platforms can be seen as covering a spectrum as illustrated in Fig. 4. One end of the spectrum consists of larger, more fault-tolerant enterprise platforms and the other end of the spectrum consists of smaller, less fault-tolerant embedded platforms.

Enterprise platforms are capable of hosting server grade machines. These platforms can generally provide a high level of node redundancy (e.g., additional computers), a high level of computational power (e.g., rack mount servers), power stability (e.g., backup generators, uninterruptable power supplies, and redundant power lines), and communication stability (e.g., multiple, high speed, and high capacity communication links). Examples of such enterprise-level platforms include Command Centers (ranging from centers such as the Pentagon to tent based data centers in the field), Navy destroyers, or airborne C2 platforms such as an Airborne Warning and Control System (AWACS).

For the sake of this paper, we define embedded platforms as those able to support only one or a few computational nodes. Such nodes may or may not be utilizing an embedded CPU architecture; are generally limited by size, weight, and power (SWaP) concerns; may have limited communication abilities and stability; and have a higher chance of platform failure. Example embedded platforms include field helicopters such as the MH-60R, unmanned vehicles such as the Predator drone, vehicle mounted targeting pods such as LITENING Pods [15], and tablets or handheld devices carried by ground forces.

In the JCAS scenario, we treat the four C2 nodes as enterprise class platforms:

- Airborne C2 – Representing a C2 air platform such as an AWACS or Joint Surveillance Target Attack Radar System (Joint STARS).
- Tactical Operations Center (TOC) or Air Support Operations Center (ASOC).
- Brigade Fire Support Cell (BDE FSC) – A clearing house that advises commanders on the use of fire support.

- Battalion FSC (BN FSC) – An FSC at the Battalion level.

The remaining platforms are treated as embedded platforms, including the following:

- JTAC – An Air Force Service member who is deployed in a forward position to direct the action of combat aircraft engaged in CAS.
- CAS AC – An aircraft providing close air support.
- Forward Air Controller (Airborne) (FAC(A)) – An Air Force Service member who directs *from the air* the action of aircraft engaged in CAS.

B. Physical Network View

The physical network view shows the networks utilized between participants, specifically the following aspects:

- The types of network resources available. In our case, there are two, (1) *radios* that provide line-of-sight communication and are generally highly constrained, and (2) *satellite communication* that is frequently highly contended.
- Which connections are available throughout a scenario (in the absence of failures), which are not, and which might or might not exist based on specifics of the deployment.

Our current descriptions focus on how line-of-sight radio communications and satellite communications connect actors in the scenarios over time. We specifically do not call out attributes such as latency or available bandwidth for a given link as this will be dependent on the hardware deployed in specific instances and as such, any such description would likely be invalid more often than not.

Throughout the JCAS scenario, the JTAC has radio communications with the FAC(A), ISR assets, Airborne C2 units, and the JFO. Continuous communication with the TOC is provided via a satellite link. The aircraft providing CAS does not have direct radio communication with the JTAC until after it has approached the target area.

The satellite communication capabilities of strike aircraft vary depending on the type of aircraft and how each is equipped. While downlinks from a satellite are common, it is common for the aircraft to not have an uplink back.

The airborne C2 is most likely capable of providing an air bridge providing beyond line-of-sight communication between all actors. This network configuration is illustrated in Fig. 5.

C. Platform Roles

A key difference between specific deployments is how platforms are utilized. For example, both embedded and enterprise platforms are capable of hosting IM services, but do not necessarily need to. Where IM services are hosted will greatly affect the fault tolerance properties of a deployment. For example, hosting IM services on embedded platforms will expose them to greater threats (such as the physical threats to aircraft) and limit the resources available for replication.

Given the capabilities of the actors in the JCAS scenario, we outline a deployment where all services are in fact hosted on the enterprise level platforms. Fig. 6 illustrates a possible deployment based on this requirement.

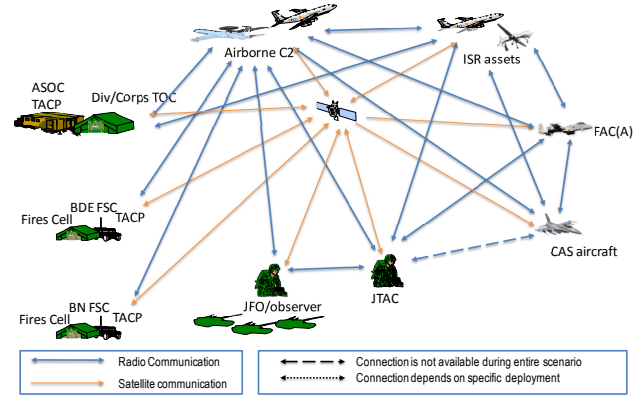


Fig. 5. JCAS Physical Network View.

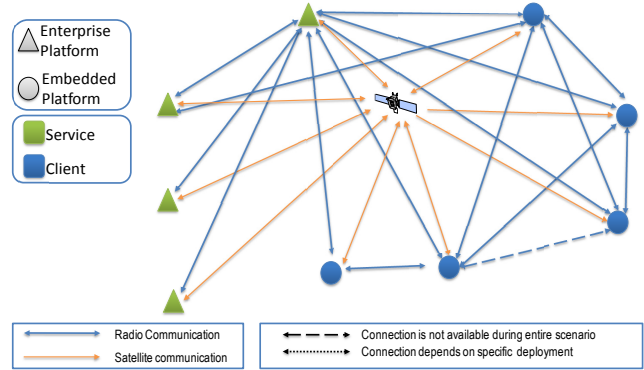


Fig. 6. Potential JCAS Deployment.

D. Logical Network View

The physical location of services and the connections between platforms do not fully describe a deployment. How platforms interact have a significant impact on how fault tolerant a given deployment is. We also describe logical network views that indicate the communication patterns across the physical links.

Consider the communication media between each enterprise platform in the JCAS Scenario. Because each platform shares the same communication bottlenecks, the Airborne C2 air bridge and the satellite link, and because each houses similar resources it is possible to collapse the enterprise platforms down into a single node on the graph. With this view a centralized deployment solution emerges, as shown in Fig. 7.

V. DEPLOYMENT FAULT TOLERANCE PROPERTIES

Based on the Information Pipeline and Service requirements presented in the scenario description, and a given deployment description, we can assess how well a set of Fault Tolerance techniques may work for a given deployment.

A. Fault Tolerance Solutions

While a detailed description of fault tolerance is outside the scope this paper, we examine here two canonical classes of fault tolerance techniques in the context of the centralized JCAS mission scenario deployment: Replication and Restart/Recovery. Replication involves running multiple copies of services and *masks* failures as long as at least one copy is up and running. Restart/recovery involves monitoring a running

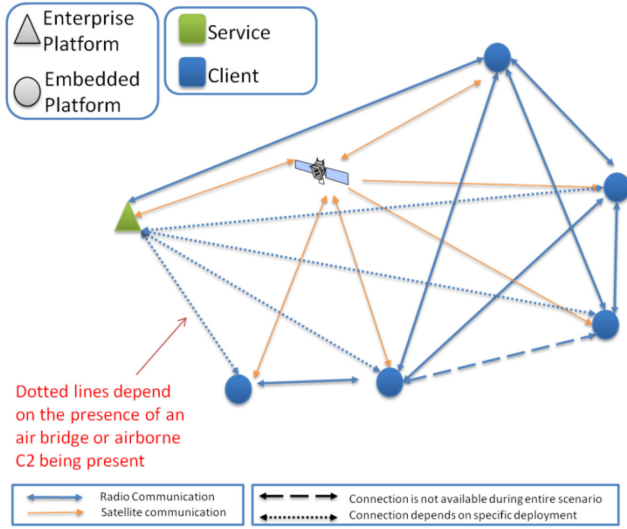


Fig. 7. Logical Network View.

service (and its process and host) and restarting the service (and process and host) if a failure occurs. During the time to detect and restart the service, the service's functionality is not available.

Replication techniques vary greatly spanning a broad spectrum of properties, e.g., required bandwidth and computational power. Our evaluations consider three general types of replication techniques:

Active replication where each replica is running and responds to every client request. This allows for seamless fail-over if one instance of a service is unavailable with the trade-off of increased computing and communication cost.

Warm passive replication where each passive replica is running but does not respond to client requests. A primary service handles the request and periodically propagates state to the passive replicas. This reduces some of the bandwidth and computational cost of active replication, but has a potentially longer fail-over time.

Cold passive replication where each passive replica is not running. A primary service logs requests and messages to a location accessible by the replicas. In the event the primary fails, the replicas start up using this logged state to ensure fault tolerance. This generally is the cheapest of the three methods with regard to computational and communication cost, but potentially has the longest fail-over time as the services need to be started up and then the logged state transferred. Replication and Restart/Recovery can each be utilized in one of three ways:

1. They can exist on the same computational node, e.g., restarting a crashed process or running multiple instances of the same service on one machine.
2. They can occur on multiple nodes, e.g., two separate computers replicating the same service or restarting a service on a different machine in the event of a crash.
3. They can occur with respect to a communication link, e.g., maintaining multiple network paths to the same destination or reinitializing a dead channel.

In our analysis, we evaluate the category of fault tolerance solutions for each deployment with respect to these three possibilities.

B. Viability of Fault Tolerance Solutions

For a fault tolerance solution to be considered viable in a given deployment there are three criteria it must meet:

1. Does a solution for each fault meet the scenario requirements of that deployment, e.g., is there a solution that recovers from a fault quickly enough to meet the scenario's timing requirements?
2. Does the deployment meet the physical requirements for a solution, e.g., can a solution be deployed on the platforms and processors available in a deployment?
3. Do the benefits of the solution outweigh the cost, e.g., does a solution impose too much overhead for the fault tolerance it provides?

Applying these criteria to the outlined JCAS scenario deployment, we derive Tables II, III, and IV describing the expected viability of the presented fault tolerance techniques.

TABLE II. JCAS DEPLOYMENT SINGLE NODE FT SOLUTIONS.

FT Technique	Suitability
Active Replication	Meets scenario requirements IM Services are housed on enterprise platforms with sufficient capabilities. Bandwidth between enterprise platforms and embedded platforms is limited and transmitting replicated data may not be desirable.
Warm Passive Replication	Solution meets scenario requirements IM Services are housed on enterprise platforms with sufficient capabilities. No unreasonable cost.
Cold Passive Replication	Solution may not meet timeliness requirements of CAS abort message IM Services are housed on enterprise platforms with sufficient capabilities. No unreasonable cost
Restart and Recovery	Solution may not meet timeliness requirements of CAS abort message Deployment meets technique requirements. No unreasonable cost

TABLE III. JCAS DEPLOYMENT MULTIPLE NODE FT SOLUTIONS.

FT Technique	Suitability
Active Replication	Meets scenario requirements Deployment meets technique requirements Bandwidth connecting enterprise platforms to embedded platforms is limited and transmitting replicated data may not be desirable.
Warm Passive Replication	Solution meets scenario requirements IM Services are housed on enterprise platforms with sufficient capabilities. No unreasonable cost
Cold Passive Replication	Solution may not meet timeliness requirements of CAS abort message IM Services are housed on enterprise platforms with sufficient capabilities. No unreasonable cost
Restart and Recovery	Solution may not meet timeliness requirements of CAS abort message Deployment meets technique requirements No unreasonable cost.

To summarize Tables II and III, either active or warm passive replication satisfies the scenario requirements for near-constant availability from the point of scenario participants. In active replication, availability would be constant and any

failures are completely masked from the participants. In warm passive, since a replica is ready to step in, failover is very rapid and should also meet the scenario requirements. In contrast, cold passive and restart/recovery incur a time during which the system is recovering and not functioning from the participants' points of view. If this recovery takes too long, it might not meet the timeliness requirements of the scenario, of which those of the *CAS abort* message are the most stringent.

Since the JCAS deployment has IM services hosted on enterprise-class servers, there should be sufficient processing resources to host the replicas needed for active and warm or cold passive replication. However, the need to send each message to multiple active replicas could overwhelm the limited links from the embedded participants, as could the performance overhead of using a group communication system. Passive replication only sends messages to a single (primary) replica, so incurs no extra cost with regard to the embedded links. Restart and recovery doesn't need to host replicas or transmit extra messages, so involves no extra resources.

TABLE IV. JCAS DEPLOYMENT COMMUNICATION LINK FT SOLUTIONS.

FT Technique	Suitability
Redundant Communication Link	Solution meets scenario requirements Only feasible if Airborne C2 (or other adequate aircraft) is able to provide an air-bridge to tactical embedded platforms No unreasonable cost
Restart Communication Link	Solution may not meet timeliness requirements of <i>CAS abort</i> message Deployment meets technique requirements No unreasonable cost

Table IV shows that having redundant communication links for failover can provide the fault masking and impression of constant availability that is needed for the most stringent timing requirements of the scenario (e.g., for the *CAS Abort* message). However, it requires sufficient infrastructure, such as a platform with multiple radios that can serve as an air bridge. Without redundant links, a communication failure could result in not meeting mission timeliness requirements, even if a failed link can be re-established.

VI. COMPARISON TO RELATED WORK

The work we presented here pulls together the disciplines of (1) the specification of mission scenarios and deployments, (2) information-centric distributed systems, and (3) mission-driven fault tolerance for distributed systems. In this section, we briefly compare our approach to existing work in each of these disciplines.

One of the most widely used frameworks for describing mission scenarios is the Department of Defense Architectural Framework (DoDAF) [23]. DoDAF describes eight different *viewpoints* for describing different aspects of a military architecture. The *Operational Viewpoint (OV)* diagram is used to describe the operations in a scenario. The *Services Viewpoint (SvcV)* and *Systems Viewpoint (SV)* describe the services, systems, and actors in a scenario and the interchanges and interconnections between them. The *Data and Information Viewpoint (DIV)* describes the data relationships in a scenario. The other viewpoints defined by DoDAF are less useful for our purposes. DoDAF was used to describe scenarios in some of

the sources that we used in our research, although many used the earlier DoDAF 1.5 version (which did not include DIV or SvcV). Other sources did not utilize a formalism, e.g., relying on text and pictures only. We included variants on the DoDAF viewpoints where possible, but incorporated only the parts that served our purposes, and supplemented them with other more technical specifications, including sequence diagrams [20]. Another specification alternative is the 4+1 Architectural View Model [14], whose *process* and *physical* views capture aspects that we capture in our scenarios. However, it does not include a specific information-centric view which is an important aspect that we need to capture.

Information-centric distributed systems rely more and more on the publish-subscribe paradigm because of its advantages in handling dynamic configurations, rich and extensible data models, and decoupling of information producers from information consumers. Many publish-subscribe messaging systems are gaining wide use, including the OMG's Data Distribution Service (DDS) [21], Java Message Service (JMS) [7], and the Advanced Message Queuing Protocol (AMQP) [25]. For this work, we are targeting a US Air Force-developed publish-subscribe system, Phoenix [6], that is service-oriented and provides not only real-time publish-subscribe, but also archival and querying and rich metadata matching, rather than simple topic-based subscriptions.

Using mission and deployment information to drive fault tolerance requirements is not a well researched area. Although the use of systems obviously influences fault tolerance requirements, e.g., spacecraft control in [2], most fault tolerance research has focused on fault tolerance of individual objects or processes. Multi-tiered fault tolerance research [13] addresses fault tolerance of communicating, interdependent services, while group communication packages, such as Spread [1] and JGroups [9], support membership and communication for groups of replicated objects. However, this work is strictly infrastructure and separate from any mission-related requirements.

VII. CONCLUSION

We have presented a reference model for documenting the properties of operational scenarios and deployments relevant to fault tolerant IM systems. We have used this model to examine a likely deployment configuration for a Joint Close Air Support scenario and example classes of fault tolerance techniques which may be utilized during such a scenario. We hypothesize that warm active replication techniques may be the most viable class of fault tolerance solutions in this deployment.

For future work we hope to validate this model by documenting additional operational scenarios and experimentally testing the viability estimates produced.

ACKNOWLEDGMENT

This work was supported by the U.S. Air Force Research Laboratory under Contract Number FA8750-10-C-0247.

REFERENCES

- [1] Y. Amir, and J. Stanton, "The Spread wide area group communication system," Technical Report CNDS 98-4, Center for Networking and Distributed Systems, Johns Hopkins University, 1998.

- [2] G. Brown, D. Bernard, and R. Rasmussen, "Attitude and articulation control for the Cassini spacecraft: a fault tolerance overview," Digital Avionics Systems Conference, November 5-9, 1995, 184-192.
- [3] R. Ceralde, "Red zones: improving the enemy ground force situation display in digital battle command and control systems," 10th International Command and Control Research & Technology Symposium (ICCRTS), McLean, VA, June 13-16, 2005.
- [4] V. Combs, R. Hillman, M. Muccio, and R. McKeel, "Joint battlespace infosphere: information management within a C2 enterprise," 10th International Command and Control Research & Technology Symposium (ICCRTS), McLean, VA, June 13-16, 2005.
- [5] M. Gillen, J. Loyall, and J. Sterling, "Dynamic quality of service management for multicast tactical communications, 14th IEEE Computer Society Symposium on Object/Component/Service-oriented Real-time Distributed Computing (ISORC), Newport Beach, CA, March 28-31, 2011.
- [6] R. Grant, V. Combs, J. Hanna, B. Lipa, and J. Reilly, "Phoenix: SOA based information management services," SPIE Conference on Defense Transformation and Net-Centric Systems, Orlando, FL, April 2009.
- [7] M. Hapner, R. Burrige, R. Sharma, J. Fialli, and K. Stout, "Java message service version 1.1," Sun Microsystems, April 12, 2002, <http://java.sun.com/products/jms/>. Accessed June 11, 2011.
- [8] Headquarters United States Air Force (AF/XOL), "Operation Anaconda, an air power perspective," February, 2005.
- [9] JGroups, <http://www.jgroups.org/>.
- [10] Joint Publication 3-09.3, "Close air support," July 2009 http://www.fas.org/irp/doddir/dod/jp3_09_3.pdf.
- [11] Joint Publication 3-09, "Joint fire support," June 2010 http://www.dtic.mil/doctrine/new_pubs/jp3_09.pdf.
- [12] Joint Publication 3-50.2, "Doctrine for joint combat search and rescue," http://www.fas.org/man/dod-101/sys/ac/docs/jp3_50_2.pdf.
- [13] B. Kemme, M. Patino-Martinez, R. Jimenez-Peris, and J. Salas, "Exactly-once interaction in a multi-tier architecture," VLDB Workshop on Design, Implementation, and Deployment of Database Replication, August 2005.
- [14] P. Kruchten, "Architectural blueprints – the 4+1 view model of software architecture," *IEEE Software* 12 (6), November 1995, 42-50.
- [15] "LITENING, advanced airborne targeting and navigation pod," Federation of American Scientists Military Analysis Network, October 28, 1999, <http://www.fas.org/man/dod-101/sys/smart/litening.htm>.
- [16] J. Loyall, J. Gossett, C. Gill, R. Schantz, J. Zinky, P. Pal, R. Shapiro, C. Rodrigues, M. Atighetchi, and D. Karr, "Comparing and contrasting adaptive middleware support in wide-area and embedded distributed object applications, 21st IEEE International Conference on Distributed Computing Systems (ICDCS-21), Phoenix, Arizona, April 16-19, 2001.
- [17] J. Loyall, R. Schantz, D. Corman, J. Paunicka, and S. Fernandez, "A distributed real-time embedded application for surveillance, detection, and tracking of time critical targets," Real-time and Embedded Technology and Applications Symposium (RTAS), San Francisco, CA, March 7-10 2005.
- [18] R. Mesic, D. Thaler, D. Ochmanek, and L. Goodson, "Courses of action for enhancing U.S. Air Force "irregular warfare" capabilities, a functional solutions analysis," RAND, 2010.
- [19] A. Newman, et al, "Time sensitive/dynamic targeting analysis techniques and results," 10th International Command and Control Research & Technology Symposium, April 2005.
- [20] Object Management Group, "OMG unified modeling language infrastructure, version 2.4.1," August 2011. <http://www.omg.org/spec/UML/2.4.1/>.
- [21] Object Management Group, "Data distribution service for real-time systems version 1.2," January 2007. formal/2007-01-01, <http://www.omg.org/spec/DDS/1.2/>.
- [22] A. Paulos, A. Sinclair, and J. Loyall, "Evaluating QoS-enabled information management service in a Navy operational context," SPIE Conference on Defense Transformation and Net-Centric Systems, Orlando, FL, April 25-29, 2011.
- [23] U.S. Department of Defense, "DoD architecture framework, version 2.0, volume 2: architectural data and models, architect's guide," May 28, 2009. http://jtc.fhu.disa.mil/jtc_dri/pdfs/dodaf_v2v2.pdf.
- [24] U.S. Department of Defense, "Executive summary of the battle of Takur Ghar," May 24, 2002.
- [25] S. Vinoski, "Advanced message queuing protocol," *IEEE Internet Computing* 10 (6): 87-89.
- [26] D. Whitcomb, "Combat search and rescue in Desert Storm," Air University Press, 2006.